SIAM Conference on Computational Science and Engineering

Short Course on the ACTS Collection:

Robust and High Performance Libraries for Computational Sciences

TAU

(Performance Analysis and Tuning)

Osni Marques

Lawrence Berkeley National Laboratory (LBNL)

oamarques@lbl.gov

TAU Performance Systems Goals

- Multi-level performance instrumentation
  - Multi-language automatic source instrumentation
- Flexible and configurable performance measurement
- Widely-ported parallel performance profiling system
  - Computer system architectures and operating systems
  - Different programming languages and compilers
- Support for multiple parallel programming paradigms
  - Multi-threading, message passing, mixed-mode, hybrid
- Support for performance mapping
- Support for object-oriented and generic programming
- Integration in complex software systems and applications

Definitions – Profiling

- Profiling
  - Recording of summary information during execution
    - inclusive, exclusive time, # calls, hardware statistics, …
  - Reflects performance behavior of program entities
    - functions, loops, basic blocks
    - user-defined "semantic" entities
  - Very good for low-cost performance assessment
  - Helps to expose performance bottlenecks and hotspots
  - Implemented through
    - sampling: periodic OS interrupts or hardware counter traps
    - instrumentation: direct insertion of measurement code

Definitions – Tracing

- Tracing
  - Recording of information about significant points (events) during program execution
    - entering/exiting code region (function, loop, block, …)
    - thread/process interactions (e.g., send/receive message)
  - Save information in event record
    - timestamp
    - CPU identifier, thread identifier
    - Event type and event-specific information
  - Event trace is a time-sequenced stream of event records
  - Can be used to reconstruct dynamic program behavior
  - Typically requires code instrumentation

TAU Performance System Architecture

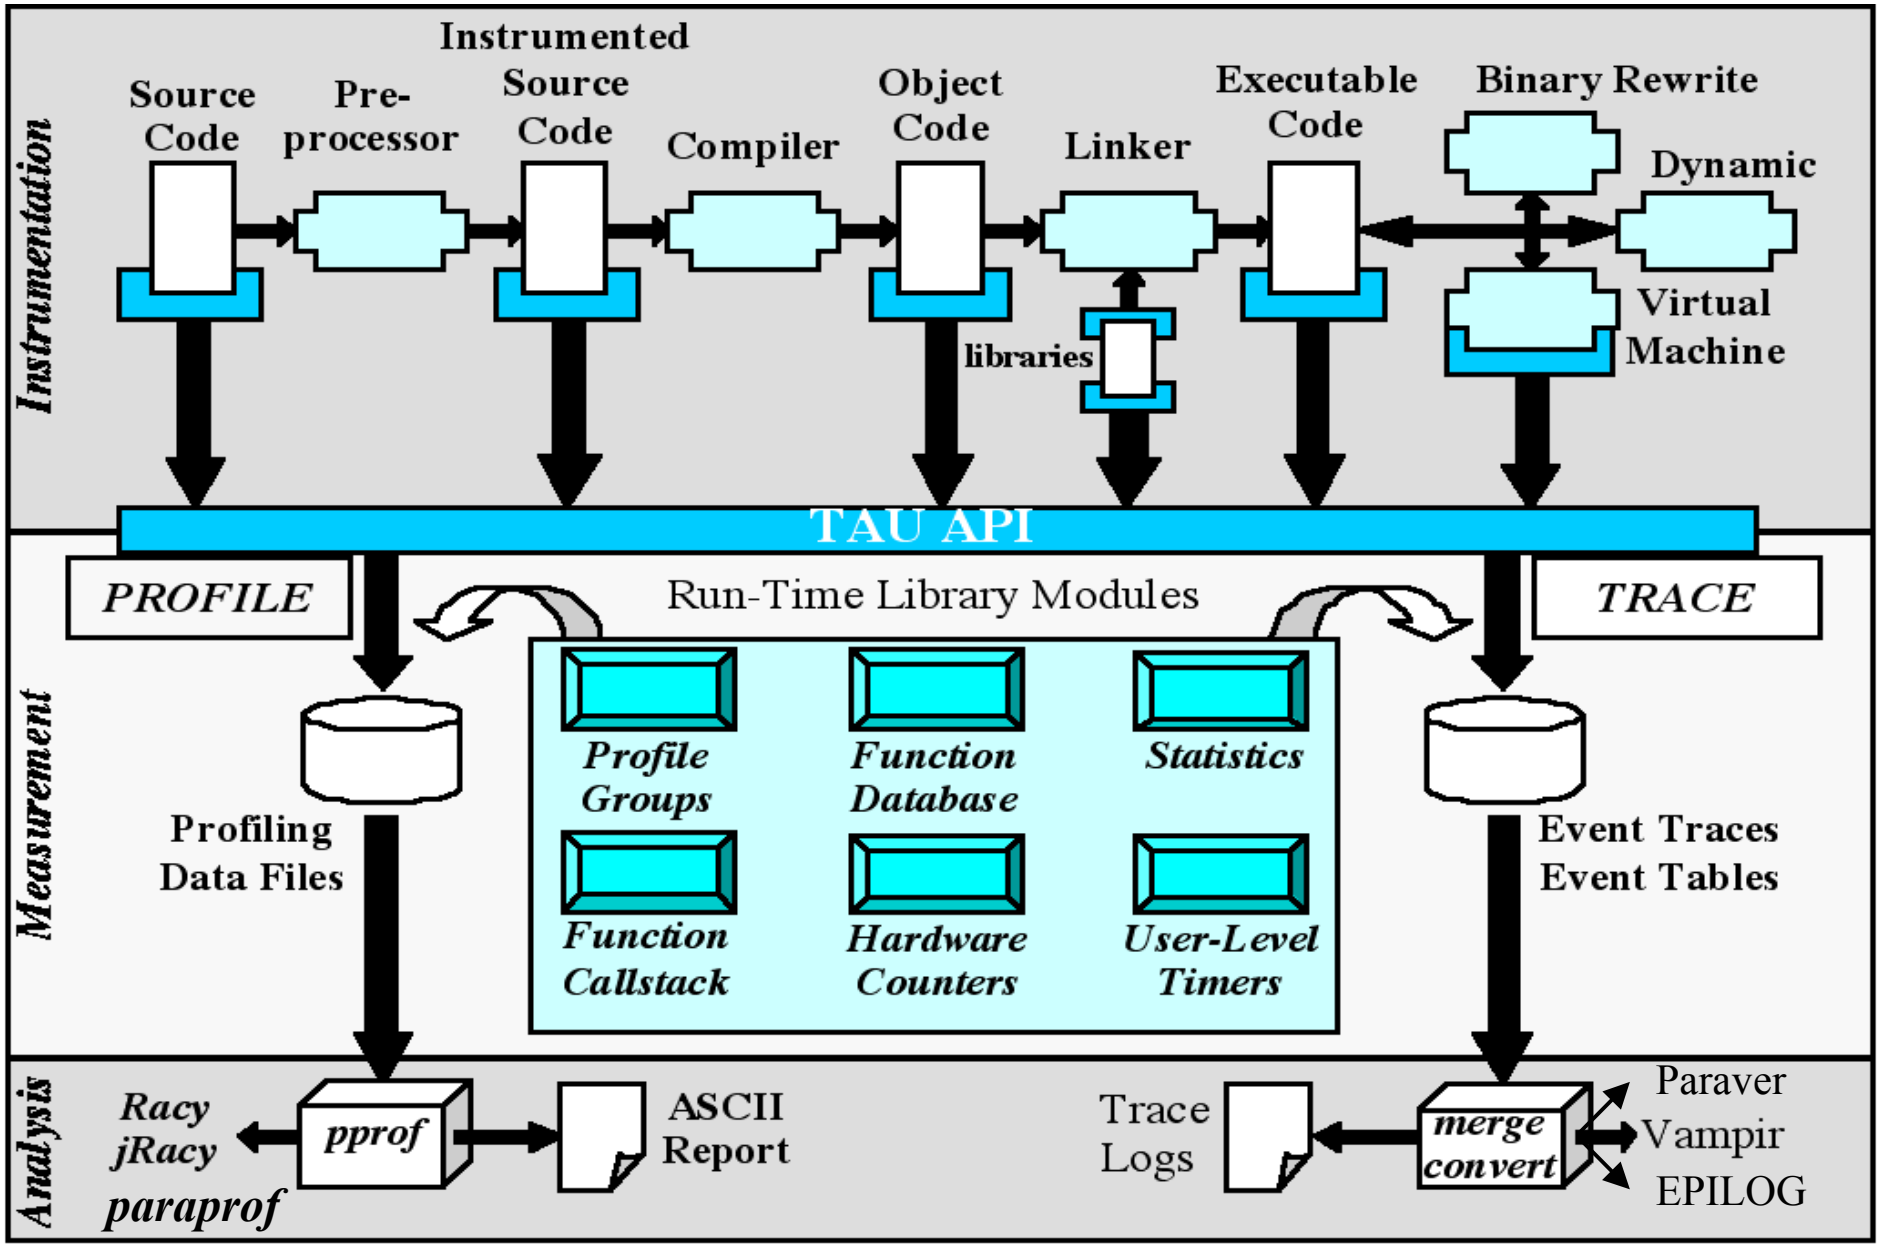

TAU Instrumentation Approach

- Support for standard program events
  - Routines
  - Classes and templates
  - Statement-level blocks
- Support for user-defined events
  - Begin/End events (“user-defined timers”)
  - Atomic events (e.g., size of memory allocated/freed)
  - Selection of event statistics
- Support definition of “semantic” entities for mapping
- Support for event groups
- Instrumentation optimization (eliminate instrumentation in lightweight routines)

TAU Instrumentation

- Flexible instrumentation mechanisms at multiple levels
  - Source code
    - manual (TAU API, TAU Component API)
    - automatic
      - C, C++, F77/90/95 (Program Database Toolkit (*PDT*))
      - OpenMP (directive rewriting with *Opari, POMP spec)*
  - Object code
    - pre-instrumented libraries (e.g., MPI using *PMPI*)
    - statically-linked and dynamically-linked
  - Executable code
    - dynamic instrumentation (pre-execution) (*DynInstAPI*)
    - virtual machine instrumentation (e.g., Java using *JVMPI*)
  - Proxy Components

Using TAU

- Install TAU

  % configure ; make clean install
- Instrument application
  - TAU Profiling API
- Typically modify application makefile
  - include TAU's stub makefile, modify variables
- Set environment variables
  - directory where profiles/traces are to be stored
  - name of merged trace file, retain intermediate trace files, etc.
- Execute application

  % mpirun –np <procs> a.out;
- Analyze performance data
  - paraprof, vampir, pprof, paraver …

TAU: *Example 1* (1/4)

Index of *http://acts.nersc.gov/tau/programs/psgesv*

Name	Last modified	Size
Parent Directory	25-May-2004 15:24	-
Makefile	25-May-2004 15:22	2k
Readme	25-May-2004 15:22	3k
psgesvdriver.f90	25-May-2004 15:22	10k
psgesvdriver.job	25-May-2004 15:22	1k

Options currently installed:

```
seaborg
[/usr/common/acts/TAU/tau-2.13.5/rs6000/lib] ls Make*
Makefile.tau-mpi-papi-pdt
Makefile.tau-mpi-papi-pdt-openmp
Makefile.tau-mpi-papi-pdt-openmp-profile-trace
Makefile.tau-mpi-papi-pdt-openmp-trace
Makefile.tau-mpi-papi-pdt-profile-trace
Makefile.tau-mpi-papi-pdt-trace
Makefile.tau-mpi-pdt
Makefile.tau-mpi-pdt-profile-trace
Makefile.tau-mpi-pdt-trace
Makefile.tau-papi-pdt
Makefile.tau-papi-pdt-openmp
Makefile.tau-papi-pdt-openmp-profile-trace
Makefile.tau-papi-pdt-openmp-trace
Makefile.tau-papi-pdt-profile-trace
Makefile.tau-papi-pdt-trace
Makefile.tau-pdt
Makefile.tau-pdt-profile-trace
Makefile.tau-pdt-trace
[/usr/common/acts/TAU/tau-2.13.5/rs6000/lib]
```

option used in the example

```
#*****************************************************************************
#*                                                                           *
#* This makefile shows how to use TAU to automatically instrument and        *
#* compile a simple Fortran program that calls the ScaLAPACK routine         *
#* PSGESV to solve a system of linear equations. It requires the             *
#* modules "tau" and "scalapack", as well as "gmake"                         *
#*                                                                           *
#*****************************************************************************
```

```
# The following defines the TAU macros
# For other options see $(TAUROOTDIR)/rs6000/lib
include $(TAUROOTDIR)/rs6000/lib/Makefile.tau-mpi-pdt-profile-trace

# Comment the following line to disable TAU
USE_TAU = 1

F90             = $(TAU_F90)
LINKER          = $(TAU_LINKER)
PDTF90PARSE     = $(PDTDIR)/$(PDTARCHDIR)/bin/f90parse
TAUINSTR        = $(TAUROOTDIR)/$(CONFIG_ARCH)/bin/tau_instrumentor
FFLAGS          = $(TAU_MPI_INCLUDE) $(TAU_F90_SUFFIX)
LIBS            = $(TAU_MPI_FLIBS) $(TAU_LIBS) $(TAU_FORTRANLIBS)
LDFLAGS         = -brtl -binitfini:poe_remote_main
```

```
ifdef USE_TAU
# Rule used for automatic instrumentation
COMP_RULE = $(PDTF90PARSE) $< $(FFLAGS); \
        $(TAUINSTR) $*.pdb $< -o $*.inst.f90 ; \
        $(F90) $(FFLAGS) -c $*.inst.f90 -o $@ ; \
        rm -f $*.pdb ;
else
# Disable TAU instrumentation
TAU_DEFS =
# Don't use TAU MPI wrapper library
TAU_MPI_LIBS = -L/usr/local/lib -lmpich -lgm
TAU_LIBS =
TAU_WRAPPER_LIB =
TAU_INCLUDE =
COMP_RULE = $(F90) $(FFLAGS) -c $< -o $@ ;
LIBS = $(TAU_MPI_LIBS)
LINKER = $(F90)
endif
```

TAU: *Example 1* (2/4)

psgesvdriver.int.f90

```
 PROGRAM PSGESVDRIVER
!
!     Example Program solving Ax=b via ScaLAPACK routine PSGESV
!
!     .. Parameters ..

!**** a bunch of things omitted for the sake of space ****

!     .. Executable Statements ..
!
!     INITIALIZE THE PROCESS GRID
!
      integer profiler(2)
      save profiler

      call TAU_PROFILE_INIT()
      call TAU_PROFILE_TIMER(profiler,'PSGESVDRIVER')
      call TAU_PROFILE_START(profiler)
      CALL SL_INIT( ICTXT, NPROW, NPCOL )
      CALL BLACS_GRIDINFO( ICTXT, NPROW, NPCOL, MYROW, MYCOL )

!**** a bunch of things omitted for the sake of space ****

      CALL PSGESV( N, NRHS, A, IA, JA, DESCA, IPIV, B, IB, JB, DESCB, &
                   INFO )

!**** a bunch of things omitted for the sake of space ****

      call TAU_PROFILE_STOP(profiler)
      STOP
      END
```

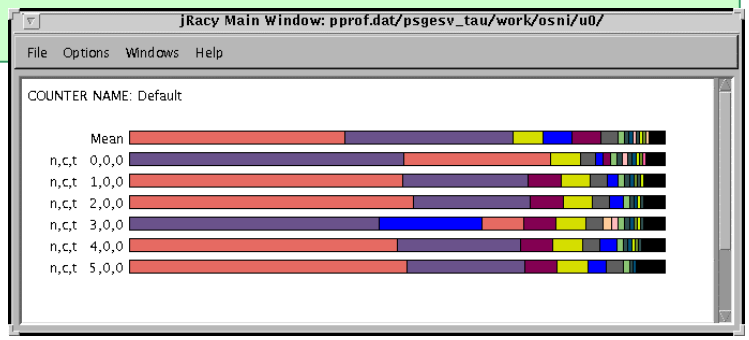

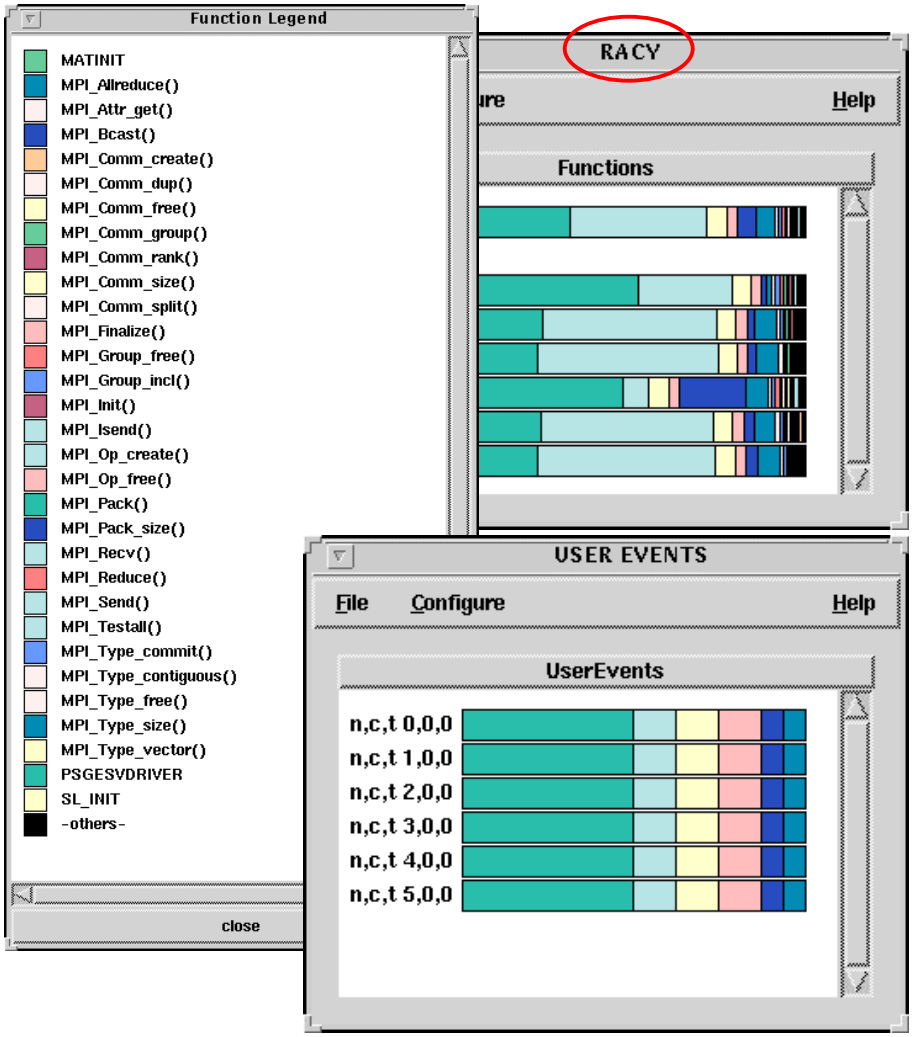

NB. ScaLAPACK routines have not been instrumented and therefore are not shown in the charts.

TAU: *Example 1* (3/4)

```
[/usr/common/homes/o/osni/work/psgesv] ls
Makefile                events.5.edf          profile.5.0.0          psgesvdriver.x
Readme                  pprof.out             psgesvdriver.err       tautrace.0.0.0.trc
events.0.edf            profile.0.0.0         psgesvdriver.f90       tautrace.1.0.0.trc
events.1.edf            profile.1.0.0         psgesvdriver.inst.f90  tautrace.2.0.0.trc
events.2.edf            profile.2.0.0         psgesvdriver.job       tautrace.3.0.0.trc
events.3.edf            profile.3.0.0         psgesvdriver.o         tautrace.4.0.0.trc
events.4.edf            profile.4.0.0         psgesvdriver.out       tautrace.5.0.0.trc
```

```
[/usr/common/homes/o/osni/work/psgesv] pprof
Reading Profile files in profile.*

NODE 0;CONTEXT 0;THREAD 0:
---------------------------------------------------------------------------------------
%Time    Exclusive    Inclusive       #Call      #Subrs  Inclusive Name
              msec   total msec                          usec/call
---------------------------------------------------------------------------------------
100.0          14           36           1         264      36477 PSGESVDRIVER
 17.3           4            6           1          16       6314 SL_INIT
 12.2           4            4          21          21        211 MPI_Bcast()
 10.7           3            3           5           5        783 MPI_Allreduce()
 10.1           3            3          16           0        231 MPI_Recv()
  2.6       0.938        0.938           1           0        938 MPI_Finalize()
  1.5       0.551        0.551           1           0        551 MPI_Comm_create()
  1.3       0.467        0.467          52           0          9 MPI_Type_commit()
  1.2       0.453        0.453           2           0        226 MPI_Comm_split()
  1.2       0.412        0.445           9           9         49 MPI_Reduce()
  1.0       0.367        0.367          50           0          7 MPI_Type_vector()
  1.0       0.364        0.364          52           0          7 MPI_Type_free()
  0.9       0.329        0.329          10           0         33 MPI_Send()
  0.6       0.216        0.216           1           0        216 MPI_Init()
  0.4       0.143        0.143           1           0        143 MPI_Comm_dup()
  0.4       0.136        0.136           7           0         19 MPI_Isend()
  0.4       0.134        0.134           8           0         17 MPI_Op_create()
  0.3       0.125        0.125          35           0          4 MPI_Type_size()
  0.2       0.065        0.065           7           0          9 MPI_Testall()
  0.1       0.043        0.043           7           0          6 MPI_Pack()
  0.1       0.034        0.034           8           0          4 MPI_Op_free()
  0.1       0.032        0.032           4           0          8 MPI_Comm_free()
  0.1       0.028        0.028           7           0          4 MPI_Pack_size()
  0.1       0.027        0.027           2           0         14 MPI_Type_contiguous()
  0.1       0.019        0.019           1           0         19 MPI_Group_incl()
  0.0       0.008        0.008           1           0          8 MPI_Group_free()
```

```
FUNCTION SUMMARY (mean):
---------------------------------------------------------------------------------------
%Time    Exclusive    Inclusive       #Call      #Subrs  Inclusive Name
              msec   total msec                          usec/call
---------------------------------------------------------------------------------------
100.0          10           36           1     222.167      36510 PSGESVDRIVER
 17.3           4            6           1          16       6329 SL_INIT
 17.0           6            6     13.3333           0        465 MPI_Recv()
 14.9           5            5          19          19        286 MPI_Bcast()
 14.5           5            5           5           5       1056 MPI_Allreduce()
  2.8           1            1           1           0       1027 MPI_Finalize()
  2.5         0.9          0.9           1           0        900 MPI_Comm_create()
  1.2       0.455        0.455           2           0        227 MPI_Comm_split()
  1.1       0.388        0.388        44.5           0          9 MPI_Type_commit()
  0.9       0.314        0.338           7           7         48 MPI_Reduce()
  0.8       0.307        0.307        44.5           0          7 MPI_Type_free()
  0.8        0.28         0.28        42.5           0          7 MPI_Type_vector()
  0.6       0.229        0.229           1           0        229 MPI_Init()
  0.6       0.214        0.214           8           0         27 MPI_Send()
  0.4       0.141        0.141           1           0        141 MPI_Comm_dup()
  0.3       0.113        0.113          31           0          4 MPI_Type_size()
  0.3       0.105        0.105           6           0         18 MPI_Op_create()
  0.3         0.1          0.1     5.33333           0         19 MPI_Isend()
  0.2       0.064        0.064     5.33333           0         12 MPI_Testall()
  0.1      0.0372       0.0372     5.33333           0          7 MPI_Pack()
  0.1       0.029        0.029           4           0          7 MPI_Comm_free()
  0.1      0.0248       0.0248           2           0         12 MPI_Type_contiguous()
  0.1      0.0232       0.0232           6           0          4 MPI_Op_free()
  0.1       0.022        0.022     5.33333           0          4 MPI_Pack_size()
  0.0      0.0167       0.0167           1           0         17 MPI_Group_incl()
  0.0      0.0107       0.0107           1           0         11 MPI_Comm_size()
  0.0     0.00817      0.00817           2           0          4 MPI_Comm_rank()
  0.0     0.00783      0.00783           1           0          8 MPI_Group_free()
  0.0      0.0075       0.0075           1           0          8 MPI_Comm_group()
  0.0      0.0075       0.0075           1           0          8 MATINIT
  0.0     0.00717      0.00717           1           0          7 MPI_Attr_get()
```

TAU: *Example 1* (4/4)

TAU: *Example 2* (1/2)

Index of *http://acts.nersc.gov/tau/programs/pdgssvx*

Name	Last modified	Size
Parent Directory	25-May-2004 15:24	-
Makefile	25-May-2004 15:22	1k
Readme	25-May-2004 16:09	2k
dcreate_matrix.c	25-May-2004 15:22	6k
lns_3937.rua	25-May-2004 15:22	559k
make.inc	25-May-2004 15:22	1k
pddrive.c	25-May-2004 15:22	5k
pddrive.job	25-May-2004 15:22	1k
pdgssvx.c	25-May-2004 15:22	43k

Makefile: *compilation rule*

```
OBJS     = pddrive.o dcreate_matrix.o pdgssvx.o
TARGET   = pddrive.x

COMPRULE = $(PDTCPARSE) $< $(COMPFLAGS) $(CDEFS) $(INCLUDEDIR); \
           $(TAUINSTR) $*.pdb $< -o $*.inst.c -g "COMPUTATION" ; \
           $(CC) $(COMPFLAGS) $(CDEFS) $(BLASDEF) $(INCLUDEDIR) -c $*.inst.c -o $@ ; \
           rm -f $*.pdb ;

$(TARGET): $(OBJS)
           $(LINKER) $(LNKFLAGS) $(OBJS) $(LIBS) $(TAULIBS) -lm -o $@

# Compilation rule
.c.o:
           $(COMPRULE)
```

```
include $(TAUROOTDIR)/rs6000/lib/Makefile.tau-mpi-papi-pdt-profile-trace
```

```
xterm
[/usr/common/homes/o/osni/work/pdgssvx] ls -la
total 28144
drwx------   2 osni     u340         8192 Jun 21 14:06 .
drwx------   4 osni     upccc        8192 Jun 21 13:27 ..
-rw-------   1 osni     u340         1428 Jun 21 13:27 Makefile
-rw-------   1 osni     u340         2401 Jun 21 13:27 Readme
-rw-------   1 osni     u340         6343 Jun 21 13:27 dcreate_matrix.c
-rw-------   1 osni     u340         6813 Jun 21 14:05 dcreate_matrix.inst.c
-rw-------   1 osni     u340         6017 Jun 21 14:05 dcreate_matrix.o
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.0.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.1.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.2.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.3.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.4.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.5.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.6.edf
-rw-------   1 osni     u340         1801 Jun 21 14:06 events.7.edf
-rw-------   1 osni     u340       572508 Jun 21 13:27 lns_3937.rua
-rw-------   1 osni     u340         1342 Jun 21 13:27 make.inc
-rw-------   1 osni     u340         5509 Jun 21 13:27 pddrive.c
-rw-------   1 osni     u340          143 Jun 21 14:06 pddrive.err
-rw-------   1 osni     u340         6178 Jun 21 14:04 pddrive.inst.c
-rwx------   1 osni     u340          412 Jun 21 13:27 pddrive.job
-rw-------   1 osni     u340         4693 Jun 21 14:05 pddrive.o
-rw-------   1 osni     u340          394 Jun 21 14:06 pddrive.out
-rwx------   1 osni     u340       936669 Jun 21 14:05 pddrive.x
-rw-------   1 osni     u340        43904 Jun 21 13:27 pdgssvx.c
-rw-------   1 osni     u340        44428 Jun 21 14:05 pdgssvx.inst.c
-rw-------   1 osni     u340        13656 Jun 21 14:05 pdgssvx.o
-rw-------   1 osni     u340         2531 Jun 21 14:06 profile.0.0.0
            1 osni     u340         2479 Jun 21 14:06 profile.1.0.0
            1 osni     u340         2492 Jun 21 14:06 profile.2.0.0
            1 osni     u340         2501 Jun 21 14:06 profile.3.0.0
            1 osni     u340         2483 Jun 21 14:06 profile.4.0.0
            1 osni     u340         2493 Jun 21 14:06 profile.5.0.0
            1 osni     u340         2489 Jun 21 14:06 profile.6.0.0
            1 osni     u340         2487 Jun 21 14:06 profile.7.0.0
            1 osni     u340      1689384 Jun 21 14:06 tautrace.0.0.0.trc
            1 osni     u340      1436496 Jun 21 14:06 tautrace.1.0.0.trc
            1 osni     u340      1698168 Jun 21 14:06 tautrace.2.0.0.trc
            1 osni     u340      1417992 Jun 21 14:06 tautrace.3.0.0.trc
            1 osni     u340      1428456 Jun 21 14:06 tautrace.4.0.0.trc
            1 osni     u340      1696488 Jun 21 14:06 tautrace.5.0.0.trc
            1 osni     u340      1441728 Jun 21 14:06 tautrace.6.0.0.trc
            1 osni     u340      1683144 Jun 21 14:06 tautrace.7.0.0.trc
sni/work/pdgssvx]
```

TAU: *Example 2* (2/2)

PAPI provides access to hardware performance counters (see *http://icl.cs.utk.edu/papi* for details and contact *acts-support@nersc.gov* for the corresponding TAU events). In this example we are just measuring FLOPS.

PARAPROF

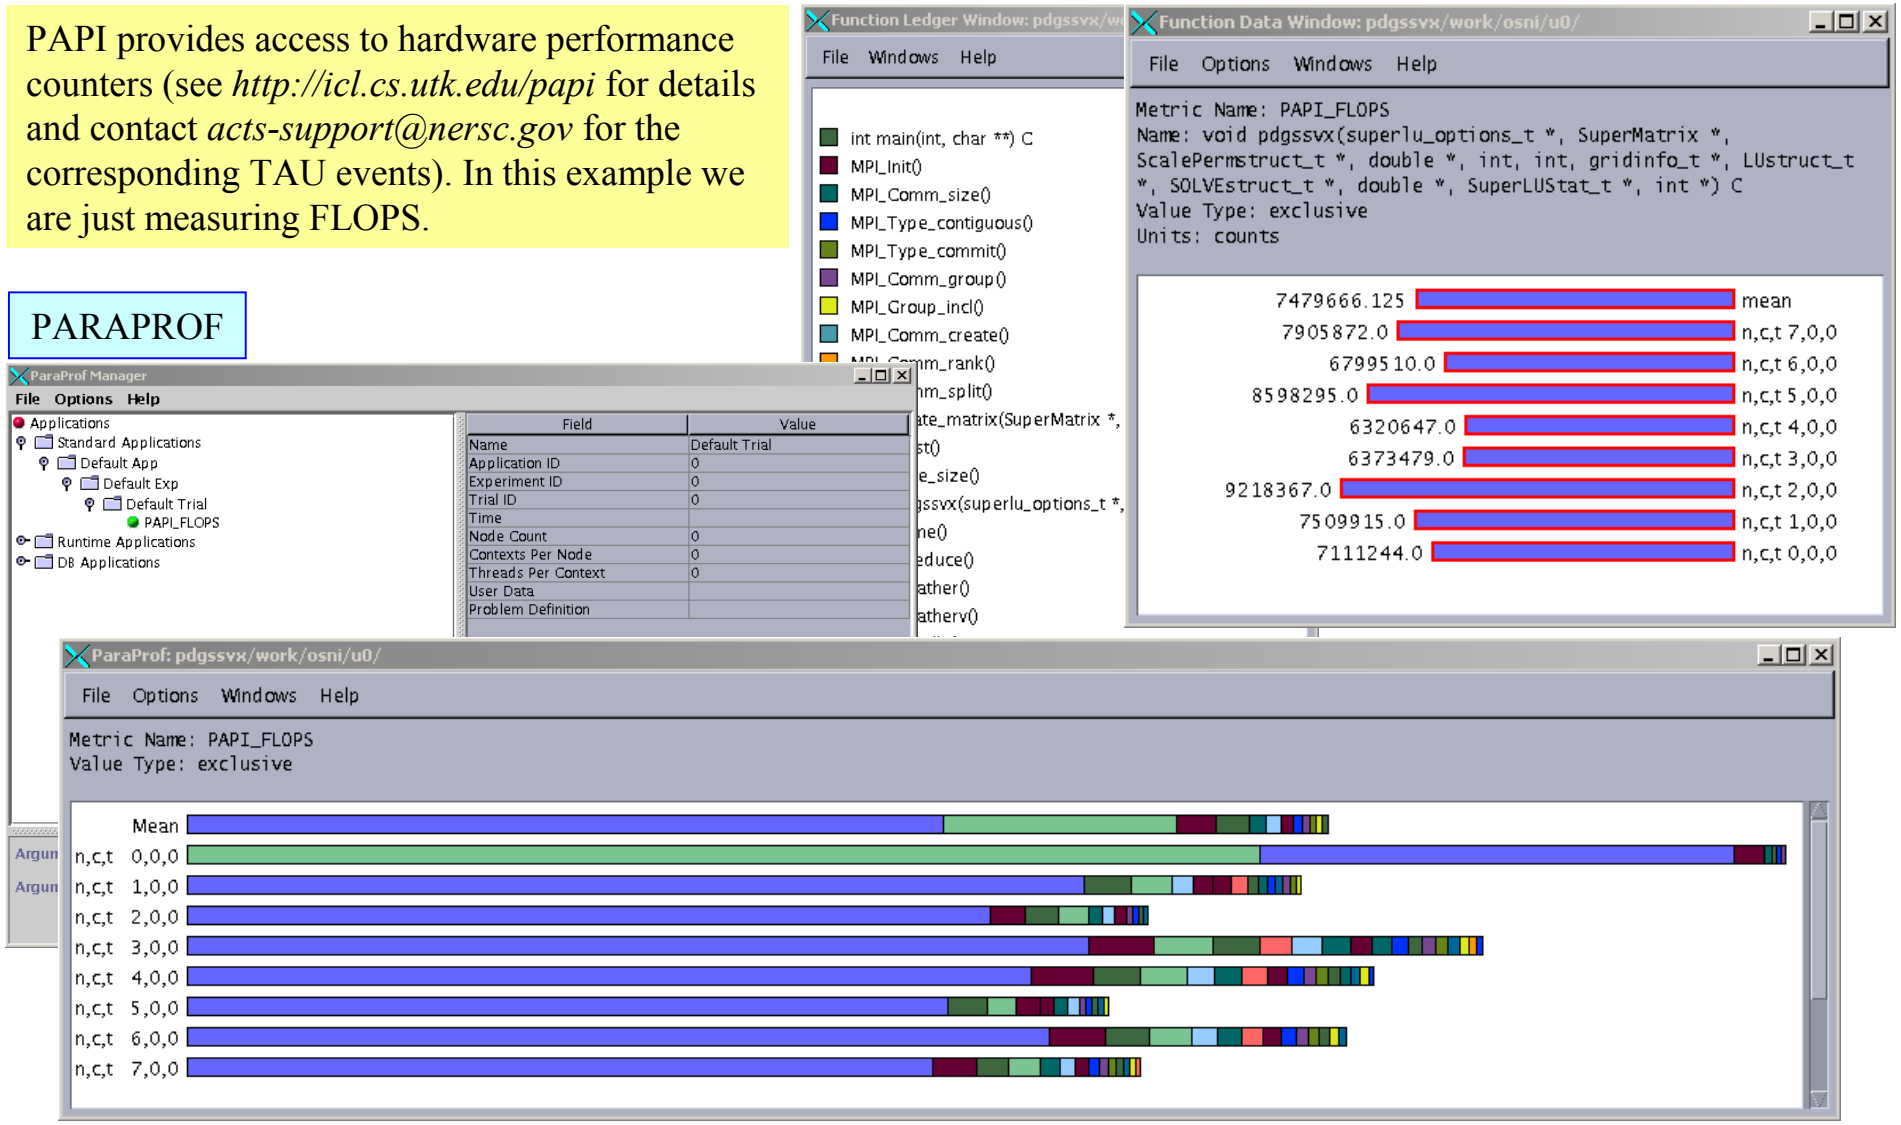

Study Case: Electronic Structure Calculations Code

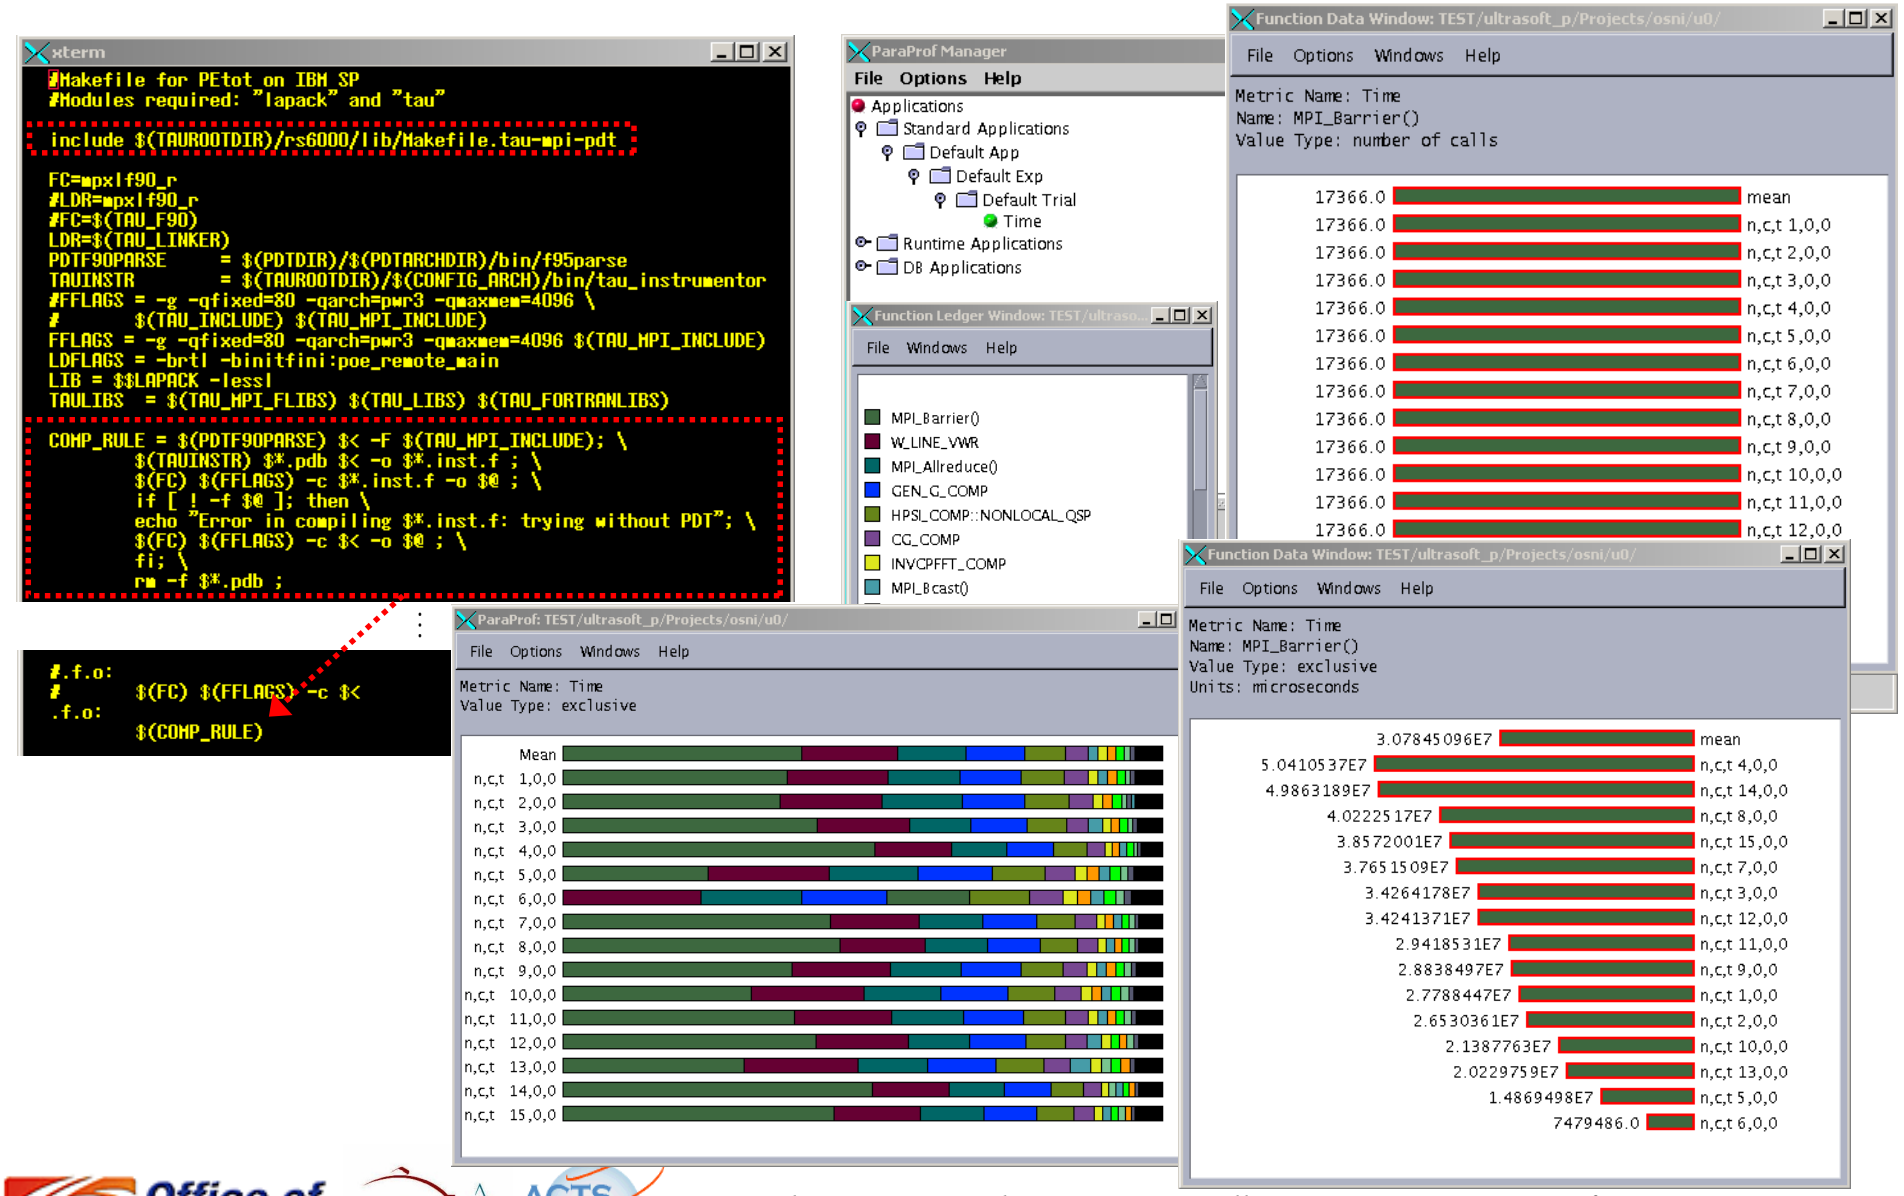

Auto Instrumentation using TAU_COMPILER

- $(TAU_COMPILER) stub Makefile variable (v2.13.7+)
- Invokes PDT parser, TAU instrumentor, compiler through **`tau_compiler.sh`** shell script
- Requires minimal changes to application Makefile
  - Compilation rules are not changed
  - User adds $(TAU_COMPILER) before compiler name
    - F90=mpxlf90
      Changes to
      F90= $(TAU_COMPILER) mpxlf90
- Passes options from TAU stub Makefile to the four compilation stages
- Uses original compilation command if an error occurs

Using TAU_COMPILER in previous example

```
#***************************************************************************
#*                                                                         *
#* This makefile shows how to use TAU to automatically instrument and      *
#* compile a simple Fortran program that calls the ScaLAPACK routine       *
#* PSGESV to solve a system of linear equations. It requires the           *
#* modules "tau" and "scalapack", as well as "gmake"                       *
#*                                                                         *
#***************************************************************************

.SUFFIXES : .f90

# The following defines the SCALAPACK macros
# include $(SCALAPACKROOTDIR)/SLmake.inc

# The following defines the TAU macros
# For other options see $(TAUROOTDIR)/rs6000/lib
include $(TAUROOTDIR)/rs6000/lib/Makefile.tau-mpi-pdt-profile-trace

F90             = $(TAU_COMPILER) $(TAU_F90)
F90SUFFIX       = -qsuffix=f=f90
LINKER          = $(TAU_LINKER)
LIBS            = $(TAU_MPI_FLIBS) $(TAU_LIBS) $(TAU_FORTRANLIBS)
LDFLAGS         = -brtl -binitfini:poe_remote_main
STLIBS          = $(SCALAPACK) $(PBLAS) $(BLACS) -lesslp2

TARGET = psgesvdriver.x
OBJS = psgesvdriver.o

$(TARGET): $(OBJS)
        $(LINKER) $(LDFLAGS) $(OBJS) -o $@ $(LIBS) $(STLIBS)

.f90.o:
        $(F90) $(F90SUFFIX) -c $<

clean:
        -@rm -f $(TARGET) *.o *.inst.f90 profile.* \
        *.trc *.edf *.pv *.pprof *.txt
```

TAU Performance System Status

- Computing platforms (selected)
  - IBM SP / pSeries, SGI Origin 2K/3K, Cray T3E / SV-1 / X1, HP (Compaq) SC (Tru64), Sun, Hitachi SR8000, NEC SX-5/6, Linux clusters (IA-32/64, Alpha, PPC, PA-RISC, Power, Opteron), Apple (G4/5, OS X), Windows
- Programming languages
  - C, C++, Fortran 77/90/95, HPF, Java, OpenMP, Python
- Thread libraries
  - pthreads, SGI sproc, Java,Windows, OpenMP
- Compilers (selected)
  - Intel KAI (KCC, KAP/Pro), PGI, GNU, Fujitsu, Sun, Microsoft, SGI, Cray, IBM (xlc, xlf), Compaq, NEC, Intel
- Full Tutorial by Sameer Shende

  http://acts.nersc.gov/events/Workshop2004/slides/tau.pdf